

In the Claims:

Please replace the claims as previously submitted with following claims in their entirety:

1. (Original) A method for dynamically partitioning a storage system cache among multiple workload classes having different quality-of-service (QoS) requirements, the cache holding data as data pages, the method comprising the steps of:

- maintaining a history of recently evicted pages for each class;
- determining a future cache size for the class based on the history and the QoS requirements, the future cache size being different than a current cache size for the class;
- determining whether the QoS requirements for the class are being met; and
- adjusting the future cache size to maximize the number of classes in which the QoS requirements are met.

2. (Original) The method as recited in claim 1, wherein the step of determining whether the QoS requirements for the class are being met includes the steps of:

- recording data concerning a QoS requirement for the class; and
- comparing the recorded data with said QoS requirement.

3. (Original) The method as recited in claim 1, wherein the step of determining a future cache size includes the steps of:

- recording cache hit data in the history of the class;
- recording the cache size corresponding to the cache hit data; and
- determining the future cache size based on the cache hit data and the respective cache sizes.

4. (Original) The method as recited in claim 1, wherein the step of adjusting the future cache size includes the steps of:

- increasing the future cache sizes of the classes whose QoS requirements are not met; and
- decreasing the future cache sizes of the classes whose QoS requirements are met.

5. (Original) The method as recited in claim 4, wherein the step of increasing the future cache sizes includes the step of setting a future cache size as a function of current cache size and the difference between cache hit data and corresponding cache sizes.

6. (Original) The method as recited in claim 4, wherein the step of decreasing the future cache sizes includes the step of setting a future cache size as a function of the current cache size, the number of classes and the difference between cache hit data and corresponding cache sizes.

7. (Original) The method as in claim 1 further comprising the step of allocating the cache space to the classes to maximize the overall cache hits if the QoS requirements for all classes are met.

8. (Original) The method as recited in claim 1, wherein the future cache size is adjusted periodically.

9. (Original) The method as recited in claim 1, wherein the future cache size is adjusted continuously on every request for data.

10. (Original) The method as recited in claim 1, wherein the the future cache size is adjusted to maximize the total class objectives.

11. (Original) A storage system capable of dynamically partitioning a system cache among multiple workload classes having different quality-of-service (QoS) requirements, the cache holding data as data pages, the system comprising:

means for maintaining a history of recently evicted pages for each class;

means for determining a future cache size for the class based on the history and the QoS requirements, the future cache size being different than a current cache size for the class;

means for determining whether the QoS requirements for the class are being met;

and

means for adjusting the future cache size to maximize the number of classes in which the QoS requirements are met.

12. (Original) The system as recited in claim 11, wherein the means for determining whether the QoS requirements for the class are being met includes:

means for recording data concerning a QoS requirement for the class; and

means for comparing the recorded data with said QoS requirement.

13. (Original) The system as recited in claim 11, wherein the means for determining a future cache size includes:

means for recording cache hit data in the history of the class;
means for recording the cache size corresponding to the cache hit data; and
means for determining the future cache size based on the cache hit data and the respective cache sizes.

14. (Original) The system as recited in claim 11, wherein the means for adjusting the future cache size includes:

means for increasing the future cache sizes of the classes whose QoS requirements are not met; and
means for decreasing the future cache sizes of the classes whose QoS requirements are met.

15. (Original) The system as recited in claim 14, wherein the means for increasing the future cache sizes includes means for setting a future cache size as a function of the current cache size and the difference between cache hit data and corresponding cache sizes.

16. (Original) The system as recited in claim 14, wherein the means for decreasing the future cache sizes includes means for setting a future cache size as a function of the current cache size, the number of classes, and the difference between cache hit data and corresponding cache sizes.

17. (Original) The system as recited in claim 11 further comprising means for allocating the cache space to the classes to maximize the overall cache hits if the QoS requirements for all classes are met.

18. (Original) The system as recited in claim 11, wherein the future cache size is adjusted periodically.

19. (Original) The system as recited in claim 11, wherein the future cache size is adjusted continuously on every request for data.

20. (Original) The system as recited in claim 11, wherein the the future cache size is adjusted to maximized the total class objectives.

21. (Currently Amended) A computer-program product for use with a storage system for dynamically partitioning a system cache among multiple workload classes having different quality-of-service (QoS) requirements, the cache holding data as data pages, the computer-program product comprising:

a computer-readable medium including instructions that are executable by a computing device;

means, provided on the computer-readable medium, for maintaining a history of recently evicted pages for each class;

means, provided on the computer-readable medium, for determining a future cache size for the class based on the history and the QoS requirements, the future cache size being different than a current cache size for the class;

means, provided on the computer-readable medium, for determining whether the QoS requirements for the class are being met; and

means, provided on the computer-readable medium, for adjusting the future cache size to maximize the number of classes in which the QoS requirements are met.

22. (Original) The computer-program product as recited in claim 21, wherein the means for determining whether the QoS requirements for the class are being met includes:

means, provided on the computer-readable medium, for recording data concerning a QoS requirement for the class; and

means, provided on the computer-readable medium, for comparing the recorded data with said QoS requirement.

23. (Original) The computer-program product as recited in claim 21, wherein the means for determining a future cache size includes:

means, provided on the computer-readable medium, for recording cache hit data in the history of the class;

means, provided on the computer-readable medium, for recording the cache size corresponding to the cache hit data; and

means, provided on the computer-readable medium, for determining the future cache size based on the cache hit data and the respective cache sizes.

24. (Original) The computer-program product as recited in claim 21,
wherein the means for adjusting the future cache size includes:

means, provided on the computer-readable medium, for increasing the future
cache sizes of the classes whose QoS requirements are not met; and

means, provided on the computer-readable medium, for decreasing the future
cache sizes of the classes whose QoS requirements are met.

25. (Original) The computer-program product as recited in claim 24,
wherein the means for increasing the future cache sizes includes means, provided on the
computer-readable medium, for setting a future cache size as a function of the current
cache size and the difference between cache hit data and corresponding cache sizes.

26. (Original) The computer-program product as recited in claim 24,
wherein the means for decreasing the future cache sizes includes means, provided on the
computer-readable medium, for setting a future cache size as a function of the current
cache size, the number of classes, and the difference between cache hit data and
corresponding cache sizes.

27. (Original) The computer-program product as recited in claim 21 further comprising means, provided on the computer-readable medium, for allocating the cache space to the classes to maximize the overall cache hits if the QoS requirements for all classes are met.

28. (Original) The computer-program product as recited in claim 21, wherein the future cache size is adjusted periodically.

29. (Original) The computer-program product as recited in claim 21, wherein the future cache size is adjusted continuously on every request for data.

30. (Currently Amended) The computer-program product as recited in claim 18, wherein the ~~the~~ future cache size is adjusted to maximize the total class objectives.